

Managing Distributed Systems with Smart Subscriptions

Robert E. Filman

Research Institute for Advanced Computer Science
NASA Ames Research Center
Mail Stop 269-1
Moffett Field, California 94305

Diana D. Lee

Science Applications International Corporation
NASA Ames Research Center
Mail Stop 269-1
Moffett Field, California 94305

Abstract *We describe an event-based, publish-and-subscribe system based on using “smart subscriptions” to recognize weakly structured events. We present a hierarchy of subscription languages (propositional, predicate, temporal and agent) of increasing expressability and computational complexity, and several algorithms (Sig, Memo, Lattice, Compile and RETE) for efficiently recognizing event matches. We have applied this system to implementing and managing distributed applications.*

Keywords: *Publish-and-subscribe, events, managing distributed applications, subscription languages, event channels, event matching algorithms.*

1 Introduction

This work arose in the context of developing a framework (the Object Infrastructure Framework, or OIF) to simplify creating distributed applications. That project developed technology to endow distributed systems with better reliability, security, quality of service and manageability by extending the standard remote-procedure-call mechanisms of distributed object technology (e.g., CORBA, Java RMI) to include discrete wrapping *injectors* on the communication paths between components. By injecting the behavior for error recovery,

behavior for error recovery, redundancy, authentication, access control, intrusion recognition, priority queue management, and so forth, OIF made substantial progress in separating and simplifying the first three concerns: reliability, security and quality of service [5].

However, much of manageability — fault diagnosis, intrusion detection, performance analysis and accounting [14] — proved resistant to a pure injector approach. Injectors provided a locus for recognizing manageability events, but not a mechanism or architecture for reporting them. This suggested extending OIF with an event mechanism. While motivated by OIF’s distributed management needs, we created the general distributed event system described in this paper. Critical issues in the design of that system were: defining what makes an event interesting and declaring to whom should an interesting event be reported.

In event systems, event *producers* asynchronously generate *events* received by *consumers*. We eschewed having producers directly send events to consumers. Such an architecture is brittle. It requires too much knowledge of the structure of a system in too many places: changes in the organization of event producers need to be reflected in every event consumer. Direct connection can also impede producers, particularly if they must interrupt critical activities to support the event-management of uncritical ones.

A solution to these problems is to position an intermediary between producers and consumers, an *event channel*. Producers funnel events to the event channel. Consumers describe to the event channel which events interest them. The event channel is responsible for forwarding appropriate events to interested consumers. This notion of event channel also goes under the rubric “publish and subscribe”: event consumers *subscribe* to (interesting) events *published* by event producers.

In this paper we consider how to code event channels so that exactly the right events are delivered to the appropriate consumers. We achieve expressiveness without overwhelming communication cost by directing only interesting events only to interested parties. Communication costs far more than processing, so it is better to expend effort checking that the communication is desired than to send volumes of uninteresting data. Of course, local processing isn’t quite free, either. We are thus addressing two issues: what *subscription languages* allow consumers to precisely describe interesting events; and which algorithms allow event channels to organize the subscription space so as to efficiently recognize events matching subscriptions.

2 Architecture

2.1 Event Channels

The two relevant interfaces for this discussion are *consumers* and *event channels*. A consumer is an object to which one can *publish* an event (encoded as a string). A consumer is entitled to do whatever it wants with that event. For example, simple consumers might print the event on a debugging screen, write the event to a log file, or update a database with some salient facet of the event. Any object can act as an event *producer* by composing a string that represents an event and invoking *publish* on some consumer with that string.

The interface *event channel* extends consumer with a *subscribe* method. *Subscribe* takes

- A reference to a consumer
- A description (a string in some *subscription language*) of the set of events interesting to that consumer
- A description of what about the existing event and environment is to be reported to the consumer (that is, the structure of an event to publish to the consumer)
- Optional *signature* information (discussed below) that can be used to optimize subscription algorithms

Subscribe returns a *ticket* for managing the subscription. With that ticket, the subscriber can modify or cancel a subscription. Event channels also include a method for obtaining the closure of the set of subscriber interests—that is, a subscription that describes the union of all the channel’s subscriptions.

Event channels, being consumers, have a *publish* method. The implementation of *publish* in an event channel considers the new event in light of the existing subscriptions (and, perhaps, past events) and publishes that event (or some derivative of the event) to every consumer whose subscription matches the event.

In OIF, every virtual address space has one global (within that address space) well-known event channel. Any application or injector that has an event to report can *publish* to that event channel. Any consumer that wishes to receive events can *subscribe* to its local event channel. Event channels on different virtual machines can *subscribe* to each other. In this way, the *publish* and *subscribe* mechanism becomes distributed, while appropriate local decisions are made about whether to distribute an event.

In OIF, local (in the same virtual machine) event channels can use not only the event itself but also information from the producer’s *environment* in matching subscriptions and composing forwarding events. OIF arranges to have the salient elements of this environment copied as part of the annotations of ordinary calls [5]. For example, a process could tag a particular call with some special symbol and recognize processes created as consequences of that call as retaining that symbol in their environment. We have used this feature to track the flow of control through a distributed system, generating events when a consequent call reaches a remote object.

OIF offered two different mechanisms for making distributed event channels become aware of each other. One was to make the application responsible for setting up the event channel network and arranging appropriate subscriptions among the nodes of that network. We are currently exploring the alternative of having information about the event channel structure be distributed by the framework “underneath” the application’s ordinary communications. This has an interesting event horizon effect: knowledge of channels interested in the effects of an action travels as fast as that action and its consequences.

2.2 Events

Event systems usually support one of two different kind of events. Most event systems define strongly-typed record-like event classes, where the class structure is globally known. In such a model, the manipulator of an event knows exactly which fields it has. Often the subscription language consists of merely describing interest in all records of a given class or subclass. This has the advantage of giving the programmer a reliable set of information on which to build—if I have an event of type *t*, I know it has fields *x*, *y*, and *z*. It has disadvantage of requiring too much commonly shared information, both in space and time. We do not want to demand that every event channel have knowledge of all kinds of events or even to posit the existence of an event definition repository. We expect the event structure to change, both as temporary event types are created to answer the questions of debugging and as new event types are created as part of the system evolution.

OIF takes the opposite tack. In OIF, events are property sets (name–value pairs), without system restrictions (or promises) as to the existence of any particular name–value pair in any particular event. OIF provides a marshaling mechanism for converting event structures to strings for transmission, and an unmarshaling mechanism for reinflating them back to property–value pairs. Thus, the string event representation:

```
“userid: Fred; time: 12:40:18; type: error;  
message: Read unhappy maknam”
```

would translate into an event object with four properties (userid, time, type and message) with the corresponding (string) values. The interpreter is responsible for doing data conversion for numeric operators. There are specific notations for strings that represent remote object references and values that are themselves events.

For debugging and system evolution, the property approach allows us to introduce new event fields into a running system. In terms of subscription languages, reference to the fields of events is straightforwardly uniform. This has the further virtue that no common understanding of event structure definitions is required across the distributed system. It has the corresponding disadvantage that we lack compile-time checks that structures will have properties not explicitly demanded in subscriptions.

3 Subscription languages

A goal of this work was to minimize uninteresting communications. Broadly, this suggests a richly expressive subscription language, where a subscriber can precisely describe which events are of interest. However, the richer the subscription language the more effort is involved both at coding time in creating the subscription interpreter and at run time in deciding if a particular subscription is satisfied by a given set of events. In OIF, we created a series of subscription languages of increasing expressiveness.¹ In OIF, we have four subscription languages: *propositional*, *predicate*, *temporal* and *agent*.

The **propositional language** deals solely with the existence of properties of events. A subscriber can express interest in A, B, and C, and any event that mentioned (as properties at the

¹ In this we are reminded of the hierarchies of automata, formal language grammars, and logics, where successively elements extend the expressibility of simpler mechanisms, often at the cost of greater computational complexity. In practice, in both formal language theory and OIF, these structures are not always strictly hierarchical.

top-level) A, B and/or C would match that subscription.

The **predicate language** provides a way to refer to the values fields of events (and subfields of contained events), constants, and values from the environment; and to combine these values with relations (e.g., “less than”) and propositional connectives (e.g., “or,” and “not”) to form a logical well-formed formula. Using a Cambridge-prefix syntax, a subscription matching error or warning events for user Joe would be:

```
(and (or (= type 'error)           (1)
      (= type 'warning))
      (= user 'Joe))
```

The **temporal language** loosens the prior restriction to single events. The propositional and predicate languages reference a single event at a time and, forward that event to the consumer. The temporal language allows for expression of relations among several events. Thus, one can talk about the existence of events E1, E2, and E3, such that E1 has occurred before E2, which occurred before E3, and which share a common user. We use RETE [7], an algorithm for rule-based forward-chaining, as our temporal matching engine. In the temporal language, subscription (1) is

```
(event (time ?t1) (userid ?u1))
(event (time ?t2) (userid ?u2))
(event (time ?t3) (userid ?u3))
(test (< ?t1 ?t2 ?t3))
(test (eq ?u1 ?u2 ?u3))
```

To deal with the finiteness of memory, we guarantee only that the most recent n events will be available for matching and that new subscriptions might recognize old events.

The **agent language** carries the implication of the Cambridge-prefix form to its logical extension. Subscriptions are themselves programs, invoked by event occurrences and able to examine the local event repository. Thus, this is a mechanism for distributing agents throughout a system and cleaning connecting a distributed agent to a local environment. Since we have not yet implemented an agent language, we have little to say about them except to note their existence at the top of the language hierarchy and

their straightforward implementation with any of the standard Lisp-like interpreters.²

In operational terms, the subscription of a subscribe method expects a string. The event channel parses this string with respect to the particular language. In our implementations, we used Cambridge prefix form as the grammatical substrate of the various subscription languages, as it is the simplest-to-parse recursive language.

4 Event channel algorithms

We developed several algorithms to improve the efficiency of recognizing matching subscriptions: *Sig*, *Memo*, *Lattice*, *Compile* and *RETE*. (We have implemented all but the fourth.)

Sig, *Memo* and *Lattice* rely on recognizing the *signature* of subscriptions. The signature of a subscription, $\mathcal{S}(s)$, is set of event properties demanded by the subscription s . For example $\mathcal{S}((\text{and (or (= type 'error) (= type 'warning)) (= user 'Joe))})$ is $\{\text{type, user}\}$. That is, any event matching this subscription must have (at least) pairs whose names are “type” and “user.” We call the properties mentioned in an event, e , the *fields* of that event, $\mathcal{F}(e)$. Both \mathcal{S} and \mathcal{F} can be represented with bit-vectors for fast subset comparisons. These algorithms are

Sig. Given event e , for each subscription, s , *Sig* checks that $\mathcal{S}(s) \subset \mathcal{F}(e)$ before evaluating s . *Sig* is a quick way of excluding certainly uninteresting events. *Sig* is appropriate for applications that generate a variety of different events and use computationally complex subscriptions.

Memo. For each unique set \mathcal{F} , *Memo* keeps a cache of those subscriptions for which $\mathcal{S}(s) \subset \mathcal{F}$. When another event with fields \mathcal{F} arrives, *Memo* needs to examine only the subscriptions in the cache. On subscription

² A similar strategy can be used for sending Java virtual machine instructions as subscriptions. The Cambridge-prefix form avoids the need for a resident compiler and makes it simpler for programs to dynamically create subscriptions at the cost of requiring a (simple) interpreter in the event-channel.

updates, Memo can either examine the power set of the signature of the changed subscription, updating the corresponding memo values, or (in practice) clear the memo table. Memo is useful for the common situation where the subscription set changes slowly and events with the same fields occur repeatedly.

Lattice extends Memo with a notion of subsumption. That is, if $\mathcal{S}(x) \subset \mathcal{S}(y)$ and $\mathcal{S}(x) \not\subset \mathcal{F}(e)$, then $\mathcal{S}(y) \not\subset \mathcal{F}(e)$, and y could not match e . In general, the signatures of subscriptions form a lattice with respect to subset. Lattice constructs the (sparse) lattice as a data structure. It works by flattening the lattice to a single path. Lattice handles subscription change more easily than Memo, and is most appropriate when there is a lot of subsumption in the subscription structure.

Compile. Each subscription can be viewed programmatically: when the subscription condition is met, perform the forwarding action. Compile treats the entire subscription set as a program by sequencing the subscriptions, and performing arbitrary compiler optimizations on the resulting program. In particular, elements such as common sub-expressions can be moved forward so as to be computed only once, tests such as $(x > 3)$ can be placed so as to shadow $(x > 7)$, and subsumptions can be realized by moving subsumed rules into the then-parts of more general subscriptions. Compile is most appropriate for a relatively static subscription set that contains a large number of common sub-expressions.

RETE. The first four algorithms deal with matching a single event to a single subscription. The temporal language matches multiple events to a subscription. In our implementation, we used the JESS [6] implementation of RETE for pattern matching.

Which is best? The optimal subscription channel algorithm is a function of the expected distribution of events and subscriptions. As we have indicated, some algorithms take advantage of an expected variety in the published events, while others do better on related or repeated event

types. Similarly, the amount of effort expended when a new subscription is received can be worthwhile only given a particular frequency of subscription changes.

5 Applications

We have implemented the event channel mechanism described here in the OIF distributed computing framework, and applied it in a demonstration application [8]. That application implemented a simulation of a distributed, competitive network management application. It used injectors to achieve quality of service (i.e., real-time performance), manageability and security. It used the event mechanism to dynamically drive “inspector” user interfaces. The event mechanism also proved critical in debugging the application, particularly as injectors could generate events on every remote invocation. Events could then be selectively scanned to get a trace of interprocess calls, and this trace could be transparently directed to both visible graphic user interfaces and textual logs.

In general, in OIF one can arbitrarily and dynamically modify the injectors of proxies or set the default behavior of a set of proxies to include a particular injector. By making an injector that generates trace events and applying that injector appropriately, the event mechanism can be made to track the patterns of interprocess calls in the system.

6 Related work

6.1 Event models.

In the taxonomy of the Barrett, Clarke, Tarr and Wise’s Framework for Event-Based Software [2], OIF’s event mechanism uses point-to-point, application-to-application communication. Modules have no explicit specification of their interfaces. OIF supports dynamic system modification and allows fully abstract naming. Our publishers are Barrett’s informers; our consumers, listeners; and our event channels, routers. The subscription mechanism effectively serves to do message transformation. We posit no de-

livery constraints beyond the underlying distributed object framework. The local event channel on each virtual machine serves as a group. Rosenblum and Wolf [11] describe a seven-component framework for event observation and notification. Within that framework, our publishers are the invoker objects and subscribers are the objects of interest. Events are explicitly generated by invoking the send event action, naming is implicit in the naming of event fields (the property-based model), observation is by explicit subscription, information is by the action of a subscription, pattern abstraction and filtering is by the pattern part of the subscription language, and the partitioning arises naturally from the set of subscriptions made. We have no explicit time model, notification is by distributed object technology calls, and the resources for sorting through subscriptions are provided by the sender and the intermediary event channels.

6.2 Event implementations.

Bates [3] argues for using a rule-based publish and subscribe system to debug heterogeneous, distributed systems. Primitive events are defined and source code is annotated so that the executing program generates event instances. Bates also uses a rule-based engine for complex event detection, fairly similar to RETE, though independently discovered.

The Elvin project is a publish-subscribe service that delivers notifications on the basis on the event's content [12]. It has an event subscription language that allows subscribers to place some constraints over the notifications, flexible definition of events that allow developers to define events as required, dynamic definition of event types, and allows the creation of new events based on old events. Elvin also introduces the idea of *quenching* that "allows event producers to receive information about what consumers are expecting of them so that they need only generate events that are in demand." In contrast to Elvin, which has a single centralized event channel, OIF's event channels are distributed.

The Ariadne Debugger in TAU stores an execution history graph of events and allows the subscriber to specify patterns using a simple

subscription language that is capable expressing temporal relations among several events but unable to express other simple prepositional or complex relations among events [13]. To compensate for the language's limitations, Ariadne "provides a scalable, spread-sheet like interface for exploring match trees."

CEDMOS is composed of event-producers and event-consumers that are connected through event-transformers. "The event transformers convert streams of incoming events into different streams of events, which are ... of interest to the event-consumers" [1]. To ease the definition the event transformers, a graphical tool facilitates the definition of complex events from simple events.

Brant and Kristensen apply events to web-based notification. Their architecture includes the notions of annotated lists, a well-worked-out datatype mechanism and a good implementation of filtering [4]. Intermetrics [10] describes a design for applying events to doing debugging of distributed, components. Luckham and Frasca apply event patterns, causal histories, filtering and aggregation to provide higher levels of abstractions for managing distributed systems [9].

7 Discussion

We have discussed the publish and subscribe mechanism in the Object Infrastructure Framework. This mechanism has proved to be a powerful tool in debugging and managing distributed systems, supporting functions such as fault diagnosis, intrusion detection, performance analysis, and accounting. Key elements of this work are the ability to inject event generators into existing components, existence within a framework that provides a continuing environmental context, the use of unstructured events, rich subscription languages, and selectable and efficient algorithms for subscription resolution. Topics for further work include (1) subscription-forwarding mechanisms that do not require tree-like branching, (2) security mechanisms for subscriptions and event channeling (including the ability of an event generator to limit who could notice his events), (3) quantifying the actual performance of different event-channel algorithms in realistic cases, (4) implementing agent sub-

scription languages, and (5) implementing subscription compilation.

Acknowledgements

The ideas expressed in this paper have emerged from the work of the MCC Object Infrastructure Project, where Dr. Filman was on assignment from Lockheed Martin Corporation and Ms. Lee an MCC employee. We have further developed these ideas at NASA Ames. We thank Stu Barrett, David Filman and Ted Linden for discussions on this subject and Cecilia Aragon, David Korsmeyer, Tarang Patel, and Alex Shaykevich for their comments on drafts of this paper.

References

- [1] Baker, D., Cassandra, A., and Rashid, M. CEDMOS: Complex Event Detection and Monitoring System. Microelectronics and Computer Technology Corporation, 1998.
- [2] Barrett, D. J., Clarke, L. A., Tarr, P. L., and Wise, A. L. An event-based software integration framework. *ACM Transactions on Software Engineering and Methodology* 5, 4 (October 1996) 378–421.
- [3] Bates, P. C. Debugging heterogeneous distributed systems using event-based models of behavior. *ACM Transactions on Computer Systems* 13, 1 (February 1995), 1–31.
- [4] Brandt, S. and Kristensen, A. Web push as an internet notification service, W3C Workshop on Push Technology, (Boston, Massachusetts, September 1997), <http://keryxsoft.hpl.hp.com/doc/ins.html>.
- [5] Filman, R. E., Barrett, S., Lee, D. D., and Linden, T. Inserting ilities by controlling communications. *CACM* (in press). <http://ic-www.arc.nasa.gov/ic/darwin/oif/leo/filman/text/oif/oif-cacm-final.pdf>
- [6] Friedman-Hill, E. J. JESS, The Java Expert System Shell. DANS98–8206 Distributed Computing Systems Sandia National Labs., Livermore, CA, (September 1998). <http://herzberg.ca.sandia.gov/jess>
- [7] Forgy, C. L. RETE: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence* 19, 1 (1982) 17–37.
- [8] Lee, D. and Filman, R. Verification of compositional software architectures. Workshop on Compositional Software Architectures, Monterey, CA, January 1998. <http://www.objs.com/workshops/ws9801/papers/paper096.doc>.
- [9] Luckham, D. C. and Frasca, B. Complex Event Processing in Distributed Systems. Stanford University Technical Report CSL-TR–98–754 (March 1998). <ftp://pavg.stanford.edu/pub/cep/fabline.ps.Z>
- [10] Ress, J. Intermetrics’ Owatch debugging technology for distributed, component-based systems. OMG–DARPA–MCC Workshop on Compositional Software Architectures (Monterey, CA, Jan. 1997). <http://www.objs.com/workshops/ws9801/papers/paper058.html>.
- [11] Rosenblum, D. S., and Wolf, A. L. A design framework for internet-scale event observation and notification. *Proceedings of the Sixth European Software Engineering Conference/ACM SIGSOFT Fifth Symposium on the Foundations of Software Engineering* (September 1997), 344–360.
- [12] Segall, B. and Arnold, D. Elvin has left the building: A publish/subscribe notification service with quenching. *Proceedings of AUUG97* (Brisbane, Queensland, Australia, September 1997).
- [13] Shende, S., Cuny, J., Hansen, L., Kundu, J., McLaughry, S., and Wolf, O. Event and state-based debugging in TAU: A Prototype. *Proceedings of SPDT’96: SIGMETRICS Symposium on Parallel and Distributed Tools* (Philadelphia, May 1996), 21–30.
- [14] Stallings, W. *SNMP, SNMP-2, and CMIP: The Practical Guide to Network-Management Standards*. Reading MA: Addison-Wesley, 1993.